

METHODS AND SYSTEMS FOR PROVIDING DYNAMIC ROUTING KEY
REGISTRATION

AN APPLICATION FOR
UNITED STATES LETTERS PATENT

By

Robby Darren Benedyk
Raleigh, North Carolina

Dan Alan Brendes
Raleigh, North Carolina

David Michael Sprague
Raleigh, North Carolina

Mark Ernest Davidson
Chapel Hill, North Carolina

Peter Joseph Marsico
Carrboro, North Carolina

BEST AVAILABLE COPY

Express Mail® mailing number: EK580268110US
Date of Deposit: 26 April 2001
I hereby certify that this paper or fee is being deposited
with the United States Postal Service "Express Mail Post
Office to Addressee" service under 37 C.F.R. 1.10 on the
date indicated above and is addressed to the
Commissioner of Patents and Trademarks, Washington,
D.C. 20231
Lillian S. Glenn



Description

METHODS AND SYSTEMS FOR PROVIDING DYNAMIC
ROUTING KEY REGISTRATION

5

Priority Application Information

This application claims the benefit of United States Provisional Patent
Application No. 60/198,967, filed April 21, 2000, the disclosure of which is
incorporated herein by reference in its entirety.

10

Technical Field

The present invention relates to the routing of signaling messages in
a converged telephony-data networking environment, and more particularly
to the automatic registration of routing key information at a gateway routing
node.

15

Background Art

The convergence of traditional telecommunication networks and
traditional data networks has given rise to a number of challenging
connectivity issues. Such connectivity issues are particularly significant in
the realm of call control signaling. More specifically, traditional public
switched telephone network (PSTN) call control signaling is performed via a
signaling system 7 (SS7) signaling protocol, while signaling within a data
network is typically performed by any of a number of signaling protocols
including: transport adapter layer interface (TALI), session initiation protocol
(SIP), session description protocol (SDP), H.323, M2UA, M3UA, SUA, etc.
In a converged communication network environment, such call control
signaling protocols are employed to provide a variety of converged or inter-
network services. These services include providing basic call setup and

teardown functionality, as well as facilitating communications-related database access. For example, call control signaling protocols are typically employed to access number portability database applications, 800 / toll-free number database applications, line information database applications, calling name database applications, home location register applications, presence service databases, telephony-to-WWW domain name servers, etc.

With regard to the call setup and teardown functionality provided by call signaling protocols, it will be appreciated that a number of switching points are typically involved in the successful completion of a call. In a traditional PSTN type network, such switching points include: end offices, tandem offices, and signal transfer points. Once again, in a pure PSTN environment, SS7 signaling messages are typically employed to facilitate such call setup operations. In a converged network, such switching points may include: end offices, softswitches, media gateway controllers, media gateways, etc. In a converged network environment, a combination of SS7 and a data network-based signaling protocol (e.g., SIP, H.323, M2UA, M3UA, etc.) may be employed to provide call setup / teardown functionality. In the case of a pure data network based communication network, the SS7 signaling protocol may be replaced completely by one or more data network signaling protocols.

As the converged network environment continues to evolve and expand, the tendency of network operators to place call switching and call service database nodes within the data network component of the converged network environment is increasing. That is to say, PSTN and wireless telephone network operators are finding the economics of data network operation favorable to the placement of signaling nodes within the data sub-network of the converged network environment, as opposed to the traditional PSTN – SS7 sub-network. As such, signaling point elements that have traditionally resided within an SS7 signaling network and been assigned unique SS7 network addresses (point codes and subsystem numbers) are now being placed within a data network, such as a TCP/IP-

based network, and consequently being assigned IP addresses and port numbers.

A detailed discussion of such data-network-based telephony nodes and associated access techniques and protocols can be found in commonly-
5 assigned, co-pending International Patent Publication No. WO 00/60812, entitled *Methods and Systems For Providing Database Node Access Control Functionality In A Communications Network Routing Node*, the disclosure of which is incorporated herein by reference in entirety.

Shown in Figure 1 is a sample converged communication network,
10 generally indicated by the numeral 100. Converged network 100 includes a signaling system 7 (SS7) network component 102 and an Internet protocol (IP) network component 104. The SS7 network component includes a service switching point (SSP) 106. The IP network component includes a
15 pair of media gateway controller (MGC) nodes 108 and 110, and a media gateway (MG) node 112. An SS7-IP signaling gateway routing node (SG) 114 connects data network nodes and SS7 network nodes. It will be appreciated that an SS7 signaling protocol is employed between SSP 106 and SG 114, while a data network signaling protocol such as TALI over TCP/IP or SCTP/IP is used to facilitate communication between SG 114 and
20 the MGC pair, 108 and 110.

It will be appreciated by those skilled in the art of SS7 communications that within an SS7 signaling network, nodes are connected via dedicated 56 kbps signaling communication links. Each signaling link provides 56 kbps of bandwidth that is dedicated to communication between
25 a pair of connected SS7 nodes. However, in an IP-based signaling network, nodes are typically connected via much faster links (typically on the order of megabits per second, depending on the underlying physical and datalink layer technologies). These high bandwidth links may be shared by a number of IP nodes simultaneously. A given path in an IP network may be shared by
30 traffic from a number of connections, which can be set up and torn down dynamically.

Because SS7 signaling links are dedicated to carrying SS7 traffic and have a fixed bandwidth, the addition of a new SS7 connection at SG 114 would require the physical installation of a new, dedicated 56 kbps SS7 signaling link. However, the addition of a new TCP/IP connection at SG 114 would simply require the sharing of existing broadband resources so as to create a new TCP/IP connection. Unlike the SS7 link creation scenario, the creation of a TCP/IP connection does not necessarily require the addition of new physical resources and, instead, can be performed dynamically via software. Consequently, the addition of and connection to a new IP based network node does not necessarily require the addition of a new physical communication link at SG 114. If existing bandwidth is sufficient, the addition of a new connection to SG 114 may only require the establishment of an additional TCP/IP connection between SG 114 and the node in the IP network with which communication is desired.

Returning now to Figure 1, it will be appreciated that an external provisioning platform 116 is adapted to communicate with the SG 114 for the purpose of provisioning signaling links, administering routing data and generally configuring services provided by the node. As such, signaling link configuration and associated routing data / routing rules must be entered or modified manually by an operator via provisioning workstation 116 each time a new MGC node is added to the network, or a change in routing preference is indicated. Such manual provisioning tasks are time intensive, costly, and prone to operator error.

Therefore, what is needed is a method and system for allowing IP connected network elements to automatically and dynamically register their presence and routing preferences at an associated network routing node, thereby minimizing or eliminating the need for manual provisioning of such configuration tasks.

Disclosure of the Invention

According to one aspect, the present invention includes a signaling gateway (SG) that is capable of providing inter-network message routing

services in a converged telephony-data network environment. The SG includes a dynamic routing key registration feature which allows Internet protocol (IP) sockets to dynamically register / de-register their routing information with the signaling gateway and subsequently direct traffic
5 towards or away from themselves without the need for manual operator intervention.

In one embodiment, an SG includes a self-registering data communication module (sDCM) that is adapted to receive and process dynamic routing key registration messages from associated IP nodes over
10 existing transmission control protocol / IP (TCP/IP) connections. Such dynamic routing key registration messages may include information that is used to register a new routing key association with a TCP/IP connection, de-register an existing routing key associated with a TCP/IP connection, or modify routing key information associated with an existing TCP/IP
15 connection. It is understood that a TCP/IP connection may be identified locally on a signaling gateway by a data structure known as a socket. Thus, in order to allow dynamic registration of routing keys associated with a TCP/IP connection, a signaling gateway may associate the routing keys with a local socket for the connection.

20 As used herein, the term "routing key" refers to a parameter or combination of parameters to be extracted from or examined in a call signaling message to determine where to route the call signaling message. Exemplary SS7 routing keys include: originating point code (OPC), destination point code (DPC), subsystem number (SSN), and circuit identifier
25 code (CIC). These routing keys have conventionally been used by SS7 nodes, such as signal transfer points to route call signaling messages to other SS7 signaling nodes. According to the present invention, IP nodes in an IP network are permitted to dynamically register SS7 routing keys in an SS7/IP signaling gateway to direct traffic to or away from themselves. This
30 dynamic registration capability in a signaling gateway node avoids the difficulties of manual registration associated with conventional routing solutions.

The sDCM card employs a dual routing key table database structure that includes both a static routing key table and a dynamic routing key table. Received dynamic routing key registration messages are used to modify information in the dynamic routing key table only. During subsequent
5 signaling message routing operations, the dynamic routing key table is searched first. The failure to locate a suitable or matching routing key entry in the dynamic routing key table results in a secondary or default search of the static routing key table.

The functions for facilitating dynamic or self-registration of IP-based
10 network elements are described herein as modules or processes. It is understood that these modules or processes may be implemented as computer-executable instructions embodied in a computer-readable medium. Alternatively, the modules or processes described herein may be implemented entirely in hardware. In yet another alternative embodiment,
15 the modules or processes described herein may be implemented as a combination of hardware and software.

The processes and modules for providing dynamic routing key registration functionality are described below as being associated with cards or subsystems within a gateway routing node. It is understood that these
20 cards or subsystems include hardware for storing and executing the processes and modules. For example, each card or subsystems described below may include one or more microprocessors, such as an x86 microprocessor available from Intel Corporation or a K series microprocessor available from AMD Corporation, and associated memory.

25 Accordingly, it is an object of the present invention to provide a routing node that facilitates dynamic self-registration by Internet protocol nodes to which it is connected.

It is another object of the present invention to provide a routing node that facilitates dynamic self-de-registration by Internet protocol nodes to
30 which it is connected.

It is yet another object of the present invention to provide a routing node that facilitates dynamic self-modification of routing key information by Internet protocol nodes to which it is connected.

5 It is yet another object of the present invention to provide a method and system for allowing IP network elements to automatically direct traffic towards or away from themselves by sending messages to a routing node.

10 It is yet another object of the present invention to provide a system and method for obtaining routing key information associated with the routing of signaling messages at a routing node that includes performing a primary lookup in a first routing key table, followed by a default lookup in second routing key table in the event that a suitable routing key entry is not located in the first routing key table.

15 It is yet another object of the present invention to provide a system and method for allowing a routing key delivered by a registration message, to override all existing, similar routing key entries in a routing key table maintained at a routing node so as to cause the routing node to direct all subsequent signaling traffic associated with the routing key to the TCP connection over which the registration message was sent.

20 Some of the objects of the invention having been stated hereinabove, other objects will become evident as the description proceeds, when taken in connection with the accompanying drawings as best described hereinbelow.

Brief Description of the Drawings

25 Figure 1 is a network diagram illustrating a converged telephony-data network.

Figure 2 is a block diagram of a conventional signaling gateway.

Figure 3 is a schematic diagram of a signaling gateway including a self-registration data communication module (sDCM) according to an embodiment of the present invention.

30 Figure 4 is a diagram of a general transport adapter layer interface (TALI) registration message structure according to an embodiment of the present invention.

Figure 5 is a table detailing TALI registration message field structures according to an embodiment of the present invention.

Figure 6 is a block diagram illustrating a TALI registration message flow through an sDCM card according to an embodiment of the present invention.

Figure 7 is a table that illustrates a sample dynamic routing key table structure according to an embodiment of the present invention.

Figure 8 is a table that illustrates a sample socket table structure according to an embodiment of the present invention.

Figure 9 is a table that illustrates sample TALI registration acknowledgment message return code values according to an embodiment of the present invention.

Figure 10 is a block diagram illustrating a TALI registration acknowledgment message flow through an sDCM card according to an embodiment of the present invention.

Figure 11 is a block diagram illustrating sDCM card response to a failed socket according to an embodiment of the present invention.

Figure 12 is a block diagram illustrating signaling message flow associated with a primary lookup in a routing key database according to an embodiment of the present invention.

Figure 13 is a block diagram illustrating signaling message flow associated with a default lookup in a routing key database according to an embodiment of the present invention.

Figure 14 is a flow chart illustrating table lookup sequences in a routing key database according to an embodiment of the present invention.

Figure 15 is a network diagram illustrating dynamic changeover functionality according to an embodiment of the present invention.

Detailed Description of the Invention

Disclosed herein are several embodiments of the present invention, all of which include a network element that performs functions similar to that of a traditional telecommunications network packet routing switch, such as a

signaling gateway (SG) routing node. Each of the embodiments described and discussed below, employs an internal architecture similar to that of high performance signal transfer point (STP) and SG products which are marketed by Tekelec as the Eagle[®] STP and IP⁷ Secure Gateway[™], respectively. A block diagram that generally illustrates the base internal architecture of the IP⁷ Secure Gateway[™] product is shown in Figure 2. A detailed description of the IP⁷ Secure Gateway[™] may be found in Tekelec publication PN/909-0767-01, Rev B, August 1999, entitled *Feature Notice IP⁷ Secure Gateway[™] Release 1.0*, the disclosure of which is incorporated by reference herein in its entirety. Similarly, a detailed description of the Eagle[®] STP may be found in the *Eagle[®] Feature Guide* PN/910-1225-01, Rev. B, January 1998, published by Tekelec, the disclosure of which is incorporated herein by reference in its entirety. The specific functional components of an IP⁷ Secure Gateway[™] for transmitting and receiving transaction capabilities application part (TCAP) messages over an Internet Protocol (IP) network are described in commonly-assigned, co-pending International Patent Publication No. WO 00/35155, the disclosure of which is incorporated herein by reference in its entirety. Similarly, the functional components of an IP⁷ Secure Gateway[™] for transmitting and receiving ISDN user part (ISUP) messages over an Internet Protocol (IP) network are described in commonly-assigned, co-pending International Patent Publication No. WO 00/35156, the disclosure of which is also incorporated herein by reference in its entirety. As described in the above referenced *Feature Notice IP⁷ Secure Gateway[™]*, an IP⁷ Secure Gateway[™] 250 includes the following subsystems: a Maintenance and Administration Subsystem (MAS) 252, a communication subsystem 254 and an application subsystem 256. MAS 252 provides maintenance communications, initial program load, peripheral services, alarm processing and system disks. Communication subsystem 254 includes an interprocessor message transport (IMT) bus that is the main communication bus among all subsystems in the IP⁷ Secure Gateway[™] 250. This high-speed

communications system functions as two 125 Mbps counter-rotating serial buses.

Application subsystem 256 includes application cards that are capable of communicating with the other cards through the IMT buses.

5 Numerous types of application cards can be incorporated into SG 250, including but not limited to: a link interface module (LIM) 258 that provides SS7 links and X.25 links, a data communication module (DCM) 260 that provides a TCP/IP interface to external nodes and an application service module (ASM) 262 that provides global title translation, gateway screening

10 and other services. A translation service module (TSM) 264 may also be provided to support triggered local number portability service. Again, it should also be appreciated that, in addition to conventional SS7 LIM cards, one or more DCM cards can be employed in a similar manner to provide for the transport of Internet Protocol (IP) encapsulated SS7 messages over an

15 IP network, as described in the above referenced *Feature Notice IP⁷ Secure GatewayTM Release 1.0* publication.

Signaling Gateway Architecture

Figure 3 illustrates a signaling gateway (SG) routing node according to an embodiment of the present invention that is generally indicated by the numeral 270. SG routing node 270 is communicatively coupled to a signaling system 7 (SS7) signaling network 280 via an SS7 signaling link 282, and to a pair of media gateway controller nodes 284 and 286 via a plurality of TCP/IP connections 288. In this simplified example, it will be

20 appreciated that the SS7 network, taken together with the TCP/IP connections, effectively constitute the functional network components of a converged telephony – data network. As further illustrated in Figure 3, SG routing node 270 includes a high-speed interprocessor message transport (IMT) communications bus 274. Communicatively coupled to IMT bus 274

25 are a number of distributed processing modules or cards including: a pair of maintenance and administration subsystem processors (MASPs) 272, an SS7 capable link Interface module (LIM) 276, and an Internet protocol (IP)

30

capable self-registration data communication module (sDCM) 278. These modules are physically connected to the IMT bus 274 such that signaling and other types of messages may be routed internally between all active cards or modules. For simplicity of illustration, only a single LIM 276 and
5 sDCM 278 are included in Figure 3. However, it should be appreciated that the distributed, multi-processor architecture of the SG routing node 270 facilitates the deployment of multiple LIM, sDCM and other cards, all of which could be simultaneously connected to and communicating via IMT bus 274.

10 MASP pair 272 implement the maintenance and administration subsystem functions described above. As the MASP pair 272 are not particularly relevant to a discussion of the flexible routing attributes of the present invention, a detailed discussion of their function is not provided herein. For a comprehensive discussion of additional MASP operations and
15 functionality, the above-referenced Tekelec IP⁷ Secure GatewayTM and Eagle[®] STP publications can be consulted.

Given the SG routing node internal architecture shown in Figure 3 and briefly discussed above, it will be appreciated that one fundamental operation of the SG 270 involves the receipt of a signaling message at LIM
20 276 from an SS7 network and the subsequent internal routing of this message to sDCM 278 for transmission via a TCP/IP communication socket to one of the pair of MGC nodes 284 or 286, and vice versa. Since the receipt and subsequent processing of SS7 message signaling units (MSUs) by a LIM card is not particularly relevant to the dynamic routing key
25 registration functionality of the present invention, a detailed discussion of such LIM operation is not provided herein. Instead, the above mentioned *Eagle[®] Feature Guide* can be consulted for a detailed discussion of LIM operation and functionality.

It should be noted that it is often the case that MGC nodes, such as
30 those shown in Figure 3, are deployed in pairs so as to provide resource redundancy. In such cases, network operators often prefer to designate one of the MGC nodes as a primary resource, while the other is held in reserve

as a backup resource. Consequently, there is no load sharing between or simultaneous operation of the two MGC nodes. When the active or primary MGC node is manually taken off-line or fails, the reserve or backup MGC node must be placed in service. The sDCM card, and more particularly, the dynamic routing key registration feature of the present invention is adapted to facilitate automatic changeover in such a scenario. As used herein, the term "changeover" refers to the process of diverting traffic from a failed signaling link to a backup signaling link. In one embodiment, a variety of transport adapter layer interface (TALI) dynamic routing key registration messages are employed to realize such self-directed MGC node behavior. It will be appreciated that other signaling protocols similar in nature to TALI (e.g., SIP, SDP, SUA, M2UA, M3UA, H.323, etc.) could also be employed to provide such functionality.

Dynamic Routing Key Registration Message Structure

Shown in Figure 4 is a sample TALI dynamic routing key registration message structure, generally indicated by the numeral 300. TALI message structure 300 includes a number of fields that are common to all TALI dynamic routing key registration messages including: a synch field 302, an opcode field 304, a length field 306, a primitive field 308, a common field 310, and a data field 312. Common field 310 further includes an operation field 314, a request / reply field 316, a success / failure code field 318.

Within a message packet, synch field 302 is used to identify the message packet as being of a transport adapter layer interface (TALI) format. As used herein "TALI" refers to the transport adapter layer interface as described in Internet Engineering Task Force (IETF) Internet Draft <draft-benedyk-sigtran-tali-01.txt> entitled "Transport Adapter Layer Interface," June 2000, the disclosure of which is incorporated herein by reference in its entirety. TALI is a protocol that defines procedures and message structures for communicating SS7 messages over a stream-oriented packet-based network, such as a TCP/IP network. However, the present invention is not limited to using TALI over TCP/IP to communicate between SS7 and IP

nodes. In an alternative embodiment of the invention, stream control transmission protocol (SCTP) over IP may be used. The stream control transmission protocol is described in RFC 2960: Stream Control Transmission Protocol, the disclosure of which is incorporated herein by
5 reference in its entirety.

Opcode field 304 identifies the type of operation associated with the message. For dynamic routing key registration related messages, an opcode value equal to "mgmt" is used. Length field 306 simply indicates the length of the message (e.g., bits, octets, etc.). Primitive field 308 is used to
10 specify a group of "mgmt" operations to which the message is applicable. A primitive field value of "rkrp" signifies a dynamic routing key registration message. RKRP operation field 314 specifies a particular operation within the group of allowed operations identified by the primitive. Message data field 312 employs a structure and contains information that are dependent
15 on the combination of opcode/primitive/operation field values (i.e., each combination could employ a different message data structure).

RKRP operation field 314 contains an integer value that is used to identify the desired "rkrp" operation. Request / reply field 316 identifies whether the "rkrp" message is a request, sent by an IP node to the SG,
20 indicating a particular type of "rkrp" action, or a reply to a previous request. Success / failure code field 318 provides a success/failure indication value as part of the reply back to an IP node for each processed request, while registration data field 312 includes specific information related to the creation, termination, or modification of a routing key - TCP/IP socket
25 association.

It will be appreciated that RKRP operation field 314, request / reply field 316, success / failure code field 318, and registration data field 312 are common to all RKRP operation related messages. The primary purpose of requiring the data structures for all RKRP operations to begin with these
30 same fields, is to provide a means for a receiver to reply to unknown RKRP messages in a consistent manner. When an sDCM card receives an RKRP request message that is not understood, the request is converted into a reply

and the success/failure code field value is used to indicate that the operation is not supported (e.g., with an RKRP reply code of 'Unsupported 'rkrp' operation, 3').

As discussed above, the specific type and quantity of information contained within a routing key registration message is a function of the character of the particular routing key with which it is associated. Shown in Figure 5 is a table 330, which provides examples of routing key types 332 and the related information or data fields that are supplied by a TALI dynamic routing key registration message. For the purposes of discussion, the routing keys shown in this example can be broadly classified as either circuit identification code (CIC) based, signaling connection control part (SCCP) based, or non-CIC / non-SCCP based. Wildcard or partial routing key descriptions are permitted, several of which are presented in table 330. It will be appreciated that wildcard routing key rules could also be defined for the CIC and SCCP based classes, as well. Such wildcard descriptions are used to facilitate default routing rules based on a partial routing key definition. Specifically, table 330 defines the data content associated with a destination point code-service indicator-originating point code (DPC-SI-OPC) routing key, a DPC-SI wildcard key, a DPC wildcard key, an SI wildcard key, and a universal default wildcard key.

As indicated in Figure 5, data fields associated with TALI dynamic routing key messages include: a set of common data fields 334 (as described above), an RKRP flag field 336, an SI field 338, a DPC field 340, an OPC field 342, a CIC range start (CICS) field 344, a CIC range end (CICE) field 346, a CIC split field 348, a new CICS (NCICS) field 350, a new CICE (NCICE) field 352, and a subsystem (SSN) field 354. Some or all of the above described data fields may be required depending upon the particular type of routing key to be registered. For example, as indicated in table 330, a registration message associated with a SCCP based routing key could include the common field values (i.e., RKRP operation field 314, request / reply field 316, success / failure code field 318), an RKRP flag value, an SI value, a DPC value, and an SSN value.

It should be noted that the RKRP flag is a 2-byte field that provides 16 possible flags that control various aspects of the dynamic routing key registration operation. In one embodiment, Bit 0 serves as an override bit that is used to control how a TCP/IP socket association for a particular routing key should be manipulated. As such, the RKRP flag determines if the dynamic routing key update transaction is intended to add a specified socket association in a "load-sharing" mode or if a new association should replace (i.e., override) all existing socket associations. It is through the use of the RKRP flag that a TCP/IP capable node, via an override-designated TCP/IP socket registration request, can re-direct and subsequently receive all traffic associated with a particular routing key.

Self-Registration Data Communication Module (sDCM) Architecture

Shown in Figure 6 is a self-registration data communication module (sDCM) of the present invention, generally indicated by reference numeral 400. sDCM 400 is connected to IMT communication bus 402 and is comprised of a number of functional processes. These processes include: a TCP/IP socket layer 404 for administering lower level TCP/IP protocol functions associated with up to 50 TCP/IP sockets. TCP/IP socket layer 404 is adapted to provide the facilities necessary to send and receive digital data over a particular physical media / physical interface, such as an Ethernet type communication link. sDCM 400 also includes a connection manager process 406 for monitoring the status of and generally managing all TCP/IP sockets, a TCP/IP socket read / write process 408 for buffering and performing basic input / output (I/O) type operations for each socket, a TALI application layer 410 for adding / removing appropriate TALI header and/or trailer information to outgoing / incoming message packets, and an SS7IPGW application layer 412 for interpreting and processing TALI messages. Of particular relevance to the present invention is a dynamic routing key process 414 which is adapted to process TALI dynamic routing key registration messages and communicate pertinent registration information to a routing database update manager process 416. Routing

database update manager process 416 is adapted to administer data table updates and generally control table lookup operations within the sDCM specific routing database, which is generally indicated by the numeral 420. In one embodiment, routing database 420 is comprised of a dynamic routing
5 key table 422, a static routing key table 424, and a socket table 426.

In the case of an outbound signaling message routing operation, it will be appreciated that routing database update manager process 416 effectively controls the sequence in which the dynamic and static table lookups occur. More particularly, the dynamic routing key table 422 is
10 always searched initially, followed by a search of the static table 424 in the event that no match is located in the dynamic data table 422.

sDCM 400 includes a message transport part (MTP) level 3 process 430 and additional functional processes beyond those shown in Figure 6. However, it will be appreciated that the MTP level 3 process and other such
15 additional functional processes are not particularly relevant to a discussion of the present invention, and are therefore not discussed in detail herein. An in depth discussion of such higher level processing functionality can be found in the above-referenced Tekelec SG and STP Feature Notice Publications.

Again, it will be appreciated that the message packets received and
20 transmitted by the sDCM card 400 may include TALI type messages, session initiation protocol (SIP), M2UA, M3UA, SUA, H.323, SCTP/IP, or other signaling protocols that may be transported via TCP/IP or similar IP based protocols. Preferred packet formats for encapsulating various types of SS7 messages in IP packets are described in the above-referenced TALI
25 IETF Internet Draft. Furthermore, functionality associated with the TALI protocol is described in commonly-assigned, co-pending International Patent Publication No. WO 00/76134, the disclosure of which is incorporated herein by reference in its entirety. Again, it will be appreciated that the concepts described in this disclosure are not dependent on the above-
30 referenced TALI signaling protocol. Other functionally similar signaling protocols are intended to be within the scope of the present invention. For example, the IETF SUA/M3UA protocol may be used.

message in a manner such that the information may be effectively used by routing database update manager **416**. In an alternate embodiment, routing database update manager process **416** may be capable of receiving a dynamic routing key message and directly processing the message.

5 In any event, routing database update manager process **416** uses the information contained within or gleaned from the dynamic routing key registration message to administer an update of dynamic routing key table **422**. Again, such dynamic routing key table update operations might include the addition of a new TCP/IP socket association, the removal of an existing
10 TCP/IP socket association, or modification of routing key information associated with an existing TCP/IP socket.

Presented in Figure 9 is a table **500** containing a sample set of return codes that are employed by an sDCM in acknowledging the receipt and subsequent processing of a dynamic routing key registration request
15 message. Each entry contained in table **500** includes a TALI return code **502**, a service indicator **504** which indicates when a return code is to be used, and a message type **506** which also determines when a return code is to be used. For example, in the event that a TALI dynamic routing key registration message is successfully received and processed by sDCM **400**,
20 a dynamic routing key registration acknowledgment message would be formulated based on the original registration message, which includes a return code value of 1 (Figure 9).

It will be appreciated that in one embodiment, a TALI dynamic routing key registration acknowledgment message is simply a copy of the received
25 dynamic routing key registration message, with the request / reply field **316** (as shown in Figure 4) set to a value indicative of a "reply", and an appropriate return code included in the success / failure code field **318** (Figure 4). It will be appreciated that in an alternate embodiment, an acknowledgment message could be constructed in a more compact format
30 so as to minimize bandwidth usage.

Shown in Figure 10 is an information flow diagram associated with a TALI dynamic routing key registration acknowledgment message. As in

previous figures, the dashed line illustrates an exemplary message flow path. Figure 10 includes sDCM card **400** as presented in Figure 6 and previously described in the preceding section. As indicated in Figure 10, routing database update manager process **416** is responsible for initiating an acknowledgment message. As discussed previously, the acknowledgment message is formulated in response to the receipt and subsequent processing of a dynamic routing key registration request message.

As such, routing database update manager process **416** directs the acknowledgment message to dynamic routing key registration process **414**, which in turn passes the message to SS7IPGW application layer **412**. SS7IPGW layer **412** determines that the message is to be transmitted via an on-card TCP/IP socket and subsequently directs the acknowledgment message to TALI application layer **410**. TALI application layer **410** appends appropriate TALI header information to the message and passes the message to the appropriate socket R/W process. In this particular example, the acknowledgment message is passed to the socket 0 R/W process **408**, and eventually transmitted to the sender of the original routing key registration message via TCP/IP socket layer **404**.

Shown in Figure 11 is an information flow diagram associated with the unanticipated or non-graceful closure of a TCP/IP connection. Once again, Figure 11 includes sDCM card **400** as presented in Figure 6 and previously described in the preceding section. In such an unanticipated connection closure scenario, an explicit dynamic routing key registration message can obviously not be communicated to sDCM **400** prior to connection failure. Instead, sDCM connection manager process **406** is responsible for monitoring the status or viability of all TCP/IP connections and subsequently notifying the routing database update manager **416** in the event of a socket failure.

It is assumed in Figure 11 that a connection has failed unexpectedly and that connection manager process **406** has observed the failure. In response, connection manager process **406** sends information regarding this connection failure to routing database update manager process **416**, which

in turn updates dynamic routing key table 422 and socket table 426 accordingly. In one embodiment, all entries in dynamic routing key table 422 associated with the failed connection are deleted, and the associated socket definition entry is also deleted from socket table 426. In an alternate
5 embodiment, all entries in dynamic routing key table 422 associated with the failed connection are left intact, and the associated socket definition entry in socket table 426 is marked with a status "unavailable."

sDCM Routing Operation

10 Shown in Figures 12 and 13 are information flow diagrams associated with the routing of a signaling message. Once again, Figures 12 and 13 include sDCM card 400 as presented in Figure 6 and previously described in the preceding section. Also, Figure 14 includes a flow chart that illustrates the basic steps associated with routing key table access on the sDCM 400,
15 and may be used in conjunction with Figures 12 and 13 to better understand routing database operation.

In the example scenario illustrated in Figure 12, it is assumed that an outbound signaling message has been sent to sDCM 400 from another communication module in a signaling gateway routing node according to an
20 embodiment of the present invention. As if, for instance, LIM 276 may internally route a signaling message to sDCM 278 via IMT bus 274, as shown in Figure 3. In any event, it will be appreciated that a signaling message is received by sDCM 400 via IMT bus 402, as indicated in Figure 12. The received signaling message requires routing instructions before
25 transmission to a destination node can be performed, and as such the routing database 420 must be accessed. As indicated in Figure 12, the signaling message is eventually received by the SS7IPGW application layer 412, which subsequently requests routing information from the routing database 420. Using information contained within the outbound signaling
30 message, one or more of the routing key tables provisioned in the routing database are accessed. More particularly, the sequence in which the dynamic and static routing key tables 422 and 424, respectively, are

accessed is a key component of the present invention. As indicated in Figure 14, dynamic routing key table 422 is accessed first. If a routing key is not found in dynamic routing key table 422 that matches the relevant information contained in the outbound signaling message, then a secondary or default routing key lookup is initiated in the static routing key table 424, as
5 generally illustrated in Figure 13.

It will be appreciated that the routing of an outbound signaling message is a complex operation and entails a number of additional steps above and beyond those discussed herein. As these additional steps are
10 not particularly relevant to the present invention, they are not explicitly presented in this disclosure. A more detailed discussion of overall signaling message routing operations may be found in the above referenced *Eagle® Feature Guide* and *Feature Notice IP⁷ Secure Gateway™* publications.

Referring to Figure 14, it will be appreciated that following receipt of
15 the outbound signaling message (ST1) from IMT bus 402, the signaling message is examined and relevant routing information is gleaned (ST2). A lookup operation is then performed in the dynamic routing key table 422 using the routing information gleaned from the signaling message (ST3), and if a routing key match is found in the dynamic routing key table 422, the
20 status of a selected TCP/IP socket is determined (ST4). It should be noted that in the event that multiple sockets are associated with the matching dynamic routing key, a specific TCP/IP socket may be selected based on a signaling link selector (SLS) parameter contained in the signaling message. In the event that the selected TCP/IP socket is available, the signaling
25 message is transmitted via the selected socket (ST5). In the event that the selected socket is not available, and there are no other available sockets associated with the matching dynamic routing key, a determination is made as to whether the destination point code associated with the destination of signaling message is accessible via a peer communication module (sDCM,
30 DCM, LIM, etc.) that is currently provisioned in the signaling gateway routing node (ST6). If such a peer communication module exists in the routing node, the signaling message is forwarded to that communication module for

routing / transmission (ST7). If such a peer communication module does not exist, the signaling message may be discarded (ST9).

In the event that the lookup in the dynamic routing key table does not yield a matching routing key entry, a secondary or default lookup operation is performed in the static routing key table 424 (ST8). If a match is found the status of a selected TCP/IP socket is determined (ST4). Again, it will be appreciated that in the event that multiple sockets are associated with the matching static routing key, a specific TCP/IP socket may be selected based on a signaling link selector (SLS) parameter contained in the signaling message. In the event that the selected TCP/IP socket is available, the signaling message is transmitted via the selected socket (ST5). In the event that the selected socket is not available, and there are no other available sockets associated with the matching static routing key, a determination is made as to whether the destination point code associated with the destination of signaling message is accessible via a peer communication module (sDCM, DCM, LIM, etc.) that is currently provisioned in the signaling gateway routing node (ST6). If such a peer communication module exists in the routing node, the signaling message is forwarded to that communication module for routing / transmission (ST7). If such a peer communication module does not exist, or if there is no routing key match found in the static routing key table 424 then the signaling message may be discarded (ST9).

Automatic Changeover

The dynamic registration procedures described herein are especially well suited to provide reliability in an IP telephony network that utilizes IP-base call control nodes, such as media gateway controllers (MGCs), to set up and tear down calls. Figure 15 is a network diagram including a pair of MGCs 284 and 286 and a signaling gateway 270. These components are the same as the correspondingly-numbered components described above with respect to Figure 3. Hence a description thereof will not be repeated herein. In the illustrated network, two stream-oriented connections 1500 and 1502 are established between MGC 284 and SG 270. Similarly, two stream-oriented connections 1504 and 1506 are established between MGC 286 and

SG 270. Stream oriented connections **1500**, **1502**, **1504**, **1506**, and **1508** may be TALI over TCP/IP connections or SCTP/IP connections. Connections **1500**, **1502**, **1504**, **1506**, and **1508** may be set up using connection establishment procedures, such as the TCP three-way
5 handshake, when MGCs **284** and **286** are brought on line.

One of the connections **1500** and **1502** may be a primary connection over which communication occurs and the other connection may be a backup connection for carrying traffic in response to failure of the first connection. Similarly, one of the connections **1504** and **1506** may be a
10 primary connection over which communication occurs and the other connection may be a backup connection for carrying call signaling traffic only in response to failure of the first connection. The present invention is not limited to two connections between communicating nodes, and it is understood that any number of primary and backup connections could be
15 used.

MGCs **284** and **286** preferably monitor the status of primary connections **1500** and **1504**. For example, MGCs **284** and **286** may determine whether the sockets associated with connections **1500** and **1504** are functioning properly. In response to detecting a failure on one of the
20 primary connections **1500** or **1504**, the MGC that manages the failed connection preferably sends a routing key registration message over the backup connection to notify sDCM **278** to start sending data over the backup connection. It would seem that this would result in two entries in dynamic routing key table **422** having the same routing keys. However, as discussed
25 above with respect to Figure 14, sDCM **278** checks the availability of a socket before sending the data over a TCP connection and if the socket indicates that the connection is unavailable, sDCM **278** looks for another socket within the routing key entry. In the automatic changeover situation, the other socket would be the socket associated with the backup connection.
30 Thus, the routing key registration procedures described herein facilitate seamless changeover when one of two connections between a signaling gateway and an IP node fail.

The same automatic changeover procedure can be used to switch communication between a primary IP node and a backup IP node. For example, MGC 284 may be a primary IP node and MGC 286 may be a backup IP node. If MGC 284 fails, MGC 286 may detect this failure using
5 inter-MGC communications and send a routing key registration request to SG 270 to direct traffic originally routed to MGC 284 to itself. It is understood that in this situation, MGC 286 would store state information of MGC 284 so that switching would occur seamlessly.

10 It will be appreciated that various details of the invention may be changed without departing from the scope of the invention. Furthermore, the foregoing description is for the purpose of illustration only, and not for the purpose of limitation--the invention being defined by the claims.